

Restricting Insider Access Through Efficient Implementation of Multi-Policy Access Control Systems

Peter Mell
peter.mell@nist.gov

James M Shook
james.shook@nist.gov

Serban Gavrilă
serban.gavrila@nist.gov

National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899

ABSTRACT

The American National Standards Institute (ANSI) has standardized an access control approach, Next Generation Access Control (NGAC), that enables simultaneous instantiation of multiple access control policies. For large complex enterprises this is critical to limiting the authorized access of insiders. However, the specifications describe the required access control capabilities but not the related algorithms. While appropriate, this leaves open the important question as to whether or not NGAC is scalable. Existing cubic reference implementations indicate that it does not. For example, the primary NGAC reference implementation took several minutes to simply display the set of files accessible to a user on a moderately sized system. To solve this problem we provide an efficient access control decision algorithm, reducing the overall complexity from cubic to linear. Our other major contribution is to provide a novel mechanism for administrators and users to review allowed access rights. We provide an interface that appears to be a simple file directory hierarchy but in reality is an automatically generated structure abstracted from the underlying access control graph that works with any set of simultaneously instantiated access control policies. Our work thus provides the first efficient implementation of NGAC while enabling user privilege review through a novel visualization approach. These capabilities help limit insider access to information (and thereby limit information leakage) by enabling the efficient simultaneous instantiation of multiple access control policies.

CCS Concepts

•Security and privacy → Access control;

Keywords

ABAC; access control; algorithms; complexity; computer security; graph theory; insider; NIST; NGAC; Policy Machine; simultaneous instantiation; XaCML

1. INTRODUCTION

Most operating systems provide simple access control mechanisms that are focused on enabling users to specify which other users have access to their files (i.e., Discretionary Access Control (DAC) [13]). However, many other access control approaches exist that provide enhanced features, especially for enterprise environments. This includes capabilities relevant to particular paradigms (e.g., Chinese Wall for conflict of interest [7] and Mandatory Access Control (MAC) for the handling of classified data [19]). Other policies provide greater simplicity in administering access control at scale (e.g., Role Based Access Control (RBAC) [1]). However, methods for protecting information with multiple models under one mechanism have been lacking, and this can result in enterprises settling for using a single simple model (e.g., DAC).

The result can be restrictions on insider access being defined very loosely; this increases the risk of insiders having unnecessary access to sensitive information and sharing that information outside of the organization. To ensure that users don't inappropriately share data, enterprises may then resort to the costly and inefficient approaches of separating different data types (e.g., military classification levels) into totally distinct and isolated networks or administering multiple independent access control systems over the same data sets. Alternately, they may accept the risk of data being leaked, which can have disastrous results (e.g., classified documents being made public).

What is needed then is an access control methodology that enables an enterprise to simultaneously implement multiple policies (e.g., DAC and MAC) to limit user access to unneeded resources. Furthermore, this methodology must be able to efficiently provide access control decisions and also enable administrators/auditors to review the access privileges on a per user/attribute basis. This ability to both make decisions and provide review are necessarily features that must be scalable to accommodate the access control policies of large enterprises.

The American National Standards Institute (ANSI) has addressed this need by standardizing an access control approach, Next Generation Access Control (NGAC) [2, 3]. The NGAC stems from and is in alignment with the Policy Machine (PM) [9], a research effort by the National Institute of Standards and Technology (NIST) to develop a general purpose Attribute Based Access Control (ABAC) framework [11]. The NGAC is designed to enable simultaneous instantiation of multiple access control policies, enabling both unified access control decisions and also review of user access capabilities. Examples for how to use NGAC to enforce the DAC, MAC, RBAC, and Chinese Wall policies are available from [8]. The NGAC specification describes what constitutes a valid implementation using set theoretic notation but does not provide implementation guidance. This approach then leaves room for multiple competing ap-

This paper is authored by an employee(s) of the United States Government and is in the public domain. Non-exclusive copying or redistribution is allowed, provided that the article citation is given and the authors and agency are clearly identified as its source.

MIST'16, October 28 2016, Vienna, Austria

ACM ISBN 978-1-4503-4571-2/16/10.

DOI: <http://dx.doi.org/10.1145/2995959.2995961>

proaches and implementations. While appropriate, this leaves open the important question as to whether or not NGAC is scalable. In this work, we find that the existing reference implementations are inefficient (using cubic and quartic algorithms) which indicates that NGAC may not be scalable. The primary NGAC reference model [14] could only scale to a test model of a couple of hundred nodes, at which point it took several minutes to visualize the set of objects available to just one user. Answering this scalability question for NGAC is critical because NGAC is the only access control methodology available promising both efficient decision and review for simultaneous instantiation of multiple access control policies.

The only other multi-policy access control methodology available is the current market leader, the eXtensible Access Control Markup Language (XACML) standard [15] from the Organization for the Advancement of Structured Information Standards (OASIS) [16]. Other related logic-based policy ABAC models (that either have no reference implementation or are not multi-policy) include $ABAC_{\alpha}$ [11], HGABAC [17], and ABAC for Web Services [20]. XACML has been shown empirically to lack scalability [18] where 3 different XACML implementations all experienced performance problems in making access control decisions where the performance decreased as the number of policies increased (each policy in XACML contains the access rules for a set of target objects). In addition, all of these logic-based policy ABAC models have been shown to be NP-complete with respect to simply determining the access attributes needed by a user to access a particular resource [5] (mapping to the satisfiability problem). Thus, these methodologies do not meet our stated need for a multi-policy system that provides for both efficient decision and review. They are then undesirable for large enterprise systems with respect to ensuring the restrictions on insider access to sensitive data to avoid information leakage by insiders.

In this work we prove that NGAC is scalable by providing linear time algorithms for both access control decisions and review of user access rights. To provide efficient decision capabilities, we took a graph theoretic view to design an efficient algorithm for access control determination. We started by transforming the NGAC set theory into a graph representation (this was straightforward as the specifications themselves often use graphs to illustrate examples). Unfortunately, the resultant graphs had unusual features and constraints (with five different types of nodes, each with its own semantics). Thus, the primary challenge was in how to apply standard graph algorithms to this representation. Our solution in general was to use breadth first search (BFS) and depth first search (DFS) variants that perform a type of topological sort as primitive operations to allow us to cascade information from one type of node to another and percolate that information through the graph until the final answers are determined. The amortized cost of the multiple searches can be shown to be linear, resulting in a linear time complexity algorithm. Furthermore, it is not linear in relation to the entire access control graph, but only to the portion of the graph relevant to a particular user. This can offer even greater speedups, avoiding the need to even traverse the entire graph.

With respect to review of user privileges, the NGAC standard does not provide any guidance on visualizing access control results. We presume the reason they do not provide this capability is because each access control policy may have its own preferred method for administrative review and user interaction. However, such a policy oriented approach isn't ideal in a system that simultaneously implements multiple policies (which is the whole point of NGAC). For example, the current NIST NGAC implementation [14] requires users to choose a particular access control policy first and then navigate just within that policy structure to review user

access (requiring the administrators and users to be knowledgeable about each access control policy and which files are covered by which policies). Because of these problems, there exists a need for a generic approach for user rights visualization that will work for any set of policies that can be instantiated within NGAC (without the staff having to understand said policies). Furthermore, this default visualization should be efficient and would ideally be automatically generated from the existing access control graph to avoid additional and excessive administrative burden.

To meet this need for review capabilities, we provide the user (or the person reviewing the user's privileges) the visual experience of traversing a typical file directory hierarchy, as used by most major operating systems. However, under the hood the user is actually traversing the NGAC access control graph. We leverage one of the graph node types (object attributes) to act as file 'directories' enabling users to access their files. The user visually sees a tree but is actually traversing a graph with an exponential number of possible paths (where we generate local views on demand to avoid exponential calculations). Since the directory tree is automatically generated from the underlying graph, it can thus provide default user access to files simultaneously protected by multiple access control policies. An interesting side effect of our approach is that there can be multiple ways for a user to access the same file, without the need to explicitly create symbolic links. Thus, a document can be both stored under a person's personal directory and under a project directory with no duplication, system inconsistency, or need to explicitly create virtual links. Our algorithms for review are based off of our algorithm for decision capabilities and thus are of linear time complexity.

In summary, the contributions of this paper include:

1. the first ever study demonstrating the scalability of the NGAC multi-policy access control system,
2. a novel visualization approach to enable review of user object access on NGAC systems, and
3. linear time algorithms for performing both access control decisions and review of user access rights.

These contributions make great strides towards enabling large scale deployments of NGAC, thus offering enterprises a powerful new capability by which to efficiently limit insider access to information (and thereby limit information leakage).

The remainder of this paper is structured as follows. Section 2 discusses related work. Section 3 provides an overview of access control graphs within NGAC and provides a definition of when a user is allowed to access an object. Section 4 presents our access control algorithms and section 5 presents our visualization approach. Section 6 concludes.

2. RELATED WORK

Of the two existing multi-policy access control methodologies, XACML is the oldest with the first version having been published in 2003. Compared to the relatively young NGAC standard (published in 2013), there exist many more implementations and it has achieved much greater adoption. Perhaps this is because XACML was available first and, up to this point, there has been a lack of compelling evidence to convince the community to use NGAC. In addition, there has been the unanswered but critical question of whether or not NGAC can scale.

Likely for these reasons, there exist only two publicly available NGAC/PM reference implementations; both are available on GitHub [10]. NIST provides a reference implementation in Java

that was the primary reference used in the development of NGAC [14]. The company Medidata provides an implementation in Ruby that they use for their software products in the medical field [12]. A third GitHub policy machine implementation is available from Colorado State, but we will not reference it further as it focuses on using the NIST PM implementation to manage application-level operating system resources in Linux environments [4].

For the NIST implementation, we evaluated version 1.5. Their code related to determining which resources are available to a particular user is cubic, which explains the slow execution time even on small test sets. We provided our linear time algorithms to NIST PM development team and they plan to use a variant of our access control algorithm as well as our visualization approach in an upcoming software release.

For the Medidata code base, we evaluated their default implementation¹ in the file ‘\lib\policy_machine.rm’ of version 1.1.0. In this, they have a method ‘accessible_objects’ that determines which files a user can access. Our analysis shows this algorithm to have an $O(nm^2)$ cubic execution time. Their method ‘is_privilege’, to determine if a user has a specific privilege for a particular object, is also quadratic. Our algorithm is linear for both these operations. We provided them our algorithm and they plan to use it to improve their default implementation.

It appears that both implementations are inefficient due to a direct translation of the set theoretic NGAC notation into computer code.

3. ACCESS CONTROL GRAPH

Using the NGAC specification [3] set theoretic definitions, we can form access control graphs as follows. There are 5 types of nodes to be created: user (u), object (o), user attribute (ua), object attribute (oa), and policy class (pc). All edges are directed. Per specification all object nodes are object attribute nodes, but there could be an object attribute that is not an object. Object attribute nodes may have edges to oa and pc nodes. However, no object attribute node may point to object nodes. User nodes are sources with edges to ua nodes. User attribute nodes may have edges to ua, oa or pc nodes.² Policy class nodes are sinks. Cycles and self-loops are prohibited. User attribute to Object attributes edges are labeled with a set of one or more allowed operations (ops) (e.g., read or write). All other edges are unlabeled. All nodes must have a path to at least one pc node (without using any ua→oa edges). For complexity evaluation purposes, the number of u, o, ua, and oa nodes are unbounded. However, the number of pc nodes and the cardinality of the ops set are assumed to be small constants.

These connectivity restrictions result in several features that we can leverage. The overall graph is a directed acyclic graph (DAG) that can be divided into two DAGs: a user DAG (with u and ua nodes) and an object DAG (with o and oa nodes). The set of u nodes act as sources for the user DAG and the set of o nodes act as sources for the object DAG. The set of ua to oa edges bridge the two DAGs and this bridge is the only place where edges are labeled, with operations (ops). We refer to the set of nodes on either side of these bridging edges as border nodes. The set of pc nodes act as sinks for both DAGs. The resulting overall graph is weakly connected.

An arbitrary access control graph can now be represented as shown in figure 1. Arrows within a set represent that nodes of that




Figure 1: Diagram showing allowed edge relationships between the five different sets of NGAC node types.

type can have edges to other nodes of that type, with no cycles allowed. This means that there are no edges between nodes within the set of pc nodes (the same is true for the set of u nodes and the set of o nodes). The arrows from the set of ua nodes to the oa nodes nodes represent the bridge edges (they contain the ops labels and connect the user and object DAGs). The bridge edges are the focal point in determining user privileges (see definition 1 below).

We now discuss how to determine user privileges. The ANSI NGAC standard provides set theoretic notation to enable computation of privileges abstracted away from any particular implementation. In this work, we describe the methodology using a graph oriented approach. Our graph theoretic derivation of the ANSI NGAC set theoretic definition of how to calculate access control is as follows³:

DEFINITION 1. For a user, u_1 , to perform an operation, op_1 , on some object, o_1 , there must exist a set of ua to oa edges with label op_1 such that the tail of each edge is reachable from u_1 and the head of each edge is reachable from o_1 and where the set of pc nodes reachable from the set of head nodes is a superset of the set of pc nodes reachable from o_1 .

This definition has three data collection components:

1. Identify the ‘active’ ua to oa bridge edges. This is the set of bridge edges that have label op_1 and where the tail is reachable from u_1 and the head is reachable from o_1 . These active edges are the ones that enable u_1 to possibly have privilege op_1 on o_1 .
2. Determine the set of pc nodes reachable from the head of each active bridge edge. The union of such pc nodes is the set of ‘covered’ policy classes for op_1 .
3. Determine the set of pc nodes reachable from o_1 . These are the policy classes that are ‘required’ to be covered.

For u_1 to perform op_1 on o_1 , the set of covered policy classes must be a, not necessarily proper, superset of the set of required policy classes.

Figure 2 shows an example NGAC access control graph which we will evaluate using definition 1⁴. Note that the dashed edges

³We don’t include the NGAC definitions here because they use completely different set theoretic notation that would require extensive explanation and that is available in the NGAC standard.

⁴The edge $ua_2 \rightarrow pc_1$ fulfills the requirement in the specification that all u and ua nodes have a path to a pc node (without using bridge edges). However, the edge is not used for determining user privileges and will not be discussed further.

¹They have other implementations that make calls out to various databases.

²In the NGAC specification, ua to pc edges are allowed but are not used for access control decisions.




Figure 2: Example NGAC access control graph.

represent the bridge edges that connect the user DAG to the object DAG. The edge label ‘r’ represents read access. In this figure, user u_1 can read o_1 and o_2 but not o_3 . We now discuss this in detail:

- o_1 requires pc_2 because there is a path connecting the two. This requirement is covered by the edge $ua_1 \rightarrow oa_1$ providing ‘read’ access (because ua_1 is reachable from u_1 , oa_1 is reachable from o_1 , and pc_2 is reachable from oa_1). Thus by definition 1, u_1 can read o_1 .
- o_2 requires pc_1 and pc_2 because there is a path connecting o_2 with both pc nodes. This requirement is covered by a combination of the edges $ua_2 \rightarrow oa_4$ (which covers the pc_1 requirement) and $ua_1 \rightarrow oa_1$ (which covers the pc_2 requirement). Note that these two bridge edges would not have fulfilled the requirements had they different labels. Thus by definition 1, u_1 can read o_2 .
- o_3 requires pc_1 and pc_2 because there is a path connecting o_3 with both pc nodes. Edge $ua_2 \rightarrow oa_4$ covers the pc_1 requirement for o_3 . However, there does not exist a ua to oa edge that will satisfy o_3 ’s requirement to cover pc_2 . Edge $ua_1 \rightarrow oa_1$ does not work because oa_1 is not reachable from o_3 (which is required in definition 1). Thus by definition 1, u_1 cannot read o_3 .

Throughout this paper, we will use Figure 2 as an example where an accountant, Bob, is working on the defense systems finances for a deathstar. In this context, we can interpret the nodes as shown in Table 1. Note how the user attribute nodes represent ‘teams’ to which Bob belongs (his own team⁵ and the death star personnel team). The object attribute nodes provide a kind of hierarchy for different projects. In this example Bob has access to his own data (oa_1 , oa_2 , and o_1) as well as the overall project and financial material (oa_4 , oa_5 , and o_2). However, Bob does not have access to any of the technical designs (oa_3 and o_3).

4. ACCESS CONTROL ALGORITHMS

We now provide a linear time complexity graph algorithm to answer two of the most common types of access control requests:

1. Is user, u_1 , allowed to perform operation, op_1 , on object, o_1 ?
2. What is the set of accessible objects for a user, u_1 , and what operations can u_1 perform on each object?

⁵We had to create ua_1 to represent Bob’s access rights because the NGAC specification does not allow creation of u to oa edges.

u_1	= ‘Bob’
ua_1	= ‘Bob Privileges’
ua_2	= ‘Death Star Personnel’
o_1	= ‘Tatooine Vacation’
o_2	= ‘Defense Systems Finances’
o_3	= ‘Energy Shield’
oa_1	= ‘Bob Personal’
oa_2	= ‘Bob Deathstar Files’
oa_3	= ‘Technical Designs’
oa_4	= ‘Deathstar Project’
oa_5	= ‘Defense Systems’
pc_1	= ‘Access Control System 1’
pc_2	= ‘Access Control System 2’

Table 1: Node Labels for Access Control Graph in Figure 2

Both of these determinations can be made through a slight variation on the same algorithm, which we refer to as PReview. Furthermore, PReview forms a basis for performing policy review in general. For example, PReview can be written in ‘reverse’ to identify all users that can access a given object through a particular operation in linear time complexity.

4.1 The PReview Algorithm

The PReview algorithm first isolates the problem to just the object DAG through labeling each border oa node reachable from u_1 with a set of operations (from the ops labels on the bridge edges from reachable border oa nodes). Then, the set of objects of ‘interest’ are found by performing a reverse BFS from the set of reachable border oa nodes (without traversing any bridge edges). If we are simply trying to determine if u_1 can access a particular object, we intersect this object with the set of objects of interest (forming a new set of objects of interest with cardinality 0 or 1).

The core of the algorithm is then to iterate over each object of interest and use the object DAG to determine which are accessible and what ops are available to the user. Each time we process an object, we will perform a topological sort DFS to collect data for the object. However, we will reuse information such that the amortized cost of all the DFSs is linear. Each DFS from an object of interest determines three data sets for that object: 1) a set of candidate ops, 2) the set of ‘covered’ policy classes for each distinct op, and 3) the set of ‘required’ policy classes for the object. With these three sets, definition 1 can be applied to determine which of the candidate ops can be used by the user on the object. In short, an op is available to the user on the object if the set of covered policy classes for that op is a superset of the set of required policy classes for that object.

In more detail, the algorithm is as follows:

1. BFS from u_1 to identify the set of reachable ua border nodes (do not traverse the oa nodes). For this set of ua border nodes, let the set of ‘active’ edges be the ua to oa out-edges.
2. For each ‘active’ edge, label the oa head node with the ops edge label (eliminating duplicates). At this point, each reachable border oa node is labeled with a set of operations.
3. Create a temporary node that is a successor of each reachable border oa node.
4. Perform a backwards BFS from the temporary node (traversing edges in reverse) to find the set of objects of ‘interest’. Do not traverse any bridge edges. Once done, delete the temporary node.

5. If the goal is to determine if u_1 can access a specific object, then intersect this object with the set of objects of interest to form a new set of objects of interest (this set will contain either a single node or be the empty set).
6. For each object of interest, perform a recursive topological sort DFS to find the reachable pc nodes. However, when performing each DFS, label all processed nodes with the information found (the set of reachable pc nodes) such that subsequent DFSs can take advantage of the previously computed information. Each object of interest then is labeled with its set of reachable pc nodes. These represent the ‘required’ pc nodes for each object.
7. While performing the DFSs from the previous step, perform an additional data propagation. When a reachable border oa node is labeled with its reachable pc nodes, associate those pc nodes with the operation labels from step 2. For example, use a dictionary where the key is the op and the value is the set of reachable pc nodes. Then use the normal functionality of the DFS to propagate these dictionaries up to the root of the tree (one of the objects of interest). When a node is being processed where multiple of its successors have these dictionaries, then union the dictionaries by taking the union of all keys; for the values take the union of the values for each key⁶.
8. For each object of interest, compare the set of required policy classes against the covered policy classed for each key from the dictionary propagated to the object in the preceding step. If for some op_1 , the set of covered policy classes is a superset of the set of required policy classes for that object, then the user can use privilege op_1 on the object.

We now look at the algorithm complexity. Steps 1 and 2 can be implemented to at most perform a single traversal of each ua node, edge in the user DAG, and each bridge edge. Steps 3 and 4 traverse each object DAG edge at most once. Step 5 takes constant processes each o node at most once. In steps 6 and 7 we store DFS results at each processed node such that the information can be reused by other DFSs. As a result, the set of executed DFSs is guaranteed to traverse each edge in the object DAG at most twice (and touch each oa node at most three times). In summation, each node and edge in the graph is then guaranteed to be touched/traversed at most 3 times (most much less and some not at all). This makes the algorithm linear with respect to the number of edges, $O(n + m)$.

4.2 Empirical Algorithm Results

In this section, we evaluate the scalability of our PReview algorithm versus the two available reference implementations (NIST PM and Medidata). For our experimental platform we used an Oracle VirtualBox Ubuntu virtual machine with two cores and 10 Gb of memory running on a commodity laptop. For software to encode the algorithms, we used Python 2.7 and NetworkX (a graph algorithms library). Faster execution times can be achieved through use of more efficient programming languages (e.g., C) but our goal was to evaluate relative performance of the algorithms. These results then are an upper bound on what can be achieved relative to execution time. With respect to memory, none of the algorithms used even a majority of the available memory and thus we do not report memory usage statistics.

⁶For example, if one successor has key/value pair $op_1 \rightarrow (pc_1)$ and another successor has key/value pairs $op_1 \rightarrow (pc_2)$ and $op_2 \rightarrow (pc_3)$, then the union of the dictionaries would be key/value pairs $op_1 \rightarrow (pc_1, pc_2)$ and $op_2 \rightarrow (pc_3)$.

Number of user nodes = $.1 \times n$
 Number of user attribute nodes = $.1 \times n$
 Number of object nodes = $.5 \times n$
 Number of object attribute nodes = $.3 \times n$
 Number of pc nodes = 3

Table 2: Proportion of Nodes of Each Type

For our empirical scalability study, we used the PReview variant that computes the set of accessible objects for a particular user and the set of operations available for each accessible object. For comparative purposes, we coded up the analogous algorithms from both NGAC reference implementations (the NIST PM [14] and Medidata [12]) using the same language and libraries. These implementations are discussed in section 2. The Medidata algorithm was perfectly analogous (identical inputs and outputs), however the NIST PM algorithm performed additional work not required to obtain our desired output. For example, the NIST PM algorithm outputs the oa nodes accessible to a user, not just the o nodes. To avoid unfairly penalizing the NIST PM algorithm, we included in our implementation only those parts relevant to obtaining the desired output.

To test the scalability of the algorithms, we generated access control graphs that varied in size from 1,000 to 700,000 nodes. We used the number of nodes, n , as the independent variable and then scaled all other graph features relative to n . The proportion of nodes of each type (u, ua o, ou, and pc) are shown in table 2. For edges, we calculated an Erdos-Renyi edge probability, p , used to create random graphs [6] such that the mean number of edges per node would be no more than 5. Then, for each candidate edge allowed by the NGAC specification, we used p to determine whether or not to place the candidate edge in the graph. The only exception is that we limited the length of the u to pc paths in the user DAG and o to pc paths in the object DAG to be at most 5. We did this for the user DAG by dividing the ua nodes into 4 groups labeled with consecutive integers. Edges leaving a node were only allowed to go to nodes in groups with higher numbered labels (edges within a group were not allowed). A similar operation was performed for the object DAG.

There do not exist any references that one can leverage for creating random NGAC graphs. Thus, we assigned the above parameters according to qualitative expert domain knowledge to create as realistic NGAC graphs as possible. To make sure that any particular parameter choice did not unfairly hamper one of the algorithms, we ran numerous experiments (not shown) where for a graph size of 2000 nodes, we varied the following parameters: proportion of nodes of a particular type (u, ua, o, oa, and pc), number of layers for the user and object DAG (to include turning off this feature), and the mean number of edges per node. We chose graphs of 2000 nodes for this experiment because that was the maximum size at which all three algorithms had a less than 20 s execution time. Some of these parameter changes produced no significant effect on execution time (e.g., number of pc nodes) while others produced significant changes (e.g., those related to the number of edges in the graph). The number of edges in the graph was affected by two factors: the number of candidate edges and the p variable. The number of candidate edges was changed by varying the proportion of ua and oa nodes and the number of layers. The p variable, used to calculate whether or not to instantiate a candidate edge, was changed by varying the target parameter for the mean number of edges per node. While we were able to change the execution times through parameter manipulation, the relative execution times between the three algorithms remained the same.




Figure 3: Execution time on graphs up to 10,000 nodes.

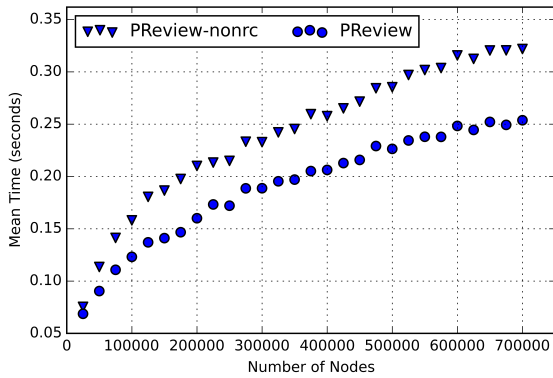


Figure 4: Execution time on graphs up to 700,000 nodes.

In the figures, we refer to our algorithm as PReview and the other two as ‘Medidata’ and ‘NIST PM’. We also provide results for an algorithm ‘PReview-nonrc’ which is simply a non-recursive version of PReview. For each data point, we took the mean of 300 trials. We limited each algorithm to taking no more than 60 s, at which point we terminated further use of that algorithm. In an actual NGAC deployment, the required response time to show a user their accessible objects is more likely to be less than 2 s.

Figure 3 shows the timing for all three primary algorithms for graphs up to 10,000 nodes. At 10,000, the PReview algorithm took a mean of .077 s to retrieve the set of objects available to a particular user. The NIST PM algorithm was 285 times slower, taking 22 s. The Medidata algorithm exceeded the 60 s limit at just 4000 nodes. Given that the required response time in an actual deployment is likely just a couple of seconds, the Medidata and NIST PM algorithms are limited to being used on graphs with less than a couple of thousand nodes (that conform to our parameters).

Figure 4 shows the performance of the PReview-nonrc and PReview algorithms on graphs up to 700,000 nodes. The PReview algorithm takes on average .254 s at 700,000 nodes. Given our assumed operation requirements of less than 2 s, this makes PReview scalable up to the largest graphs that we produced (that conform to our parameters). We did not generate larger graphs due to execution time and memory limitations on our code used to produce the NGAC graphs. Note that the non-recursive PReview-nonrc algorithm performed similarly to the recursive version but took 1.268

times longer at 700,000 nodes. It can be used in situations where the recursion is not desired or for graphs where the recursion depth might be an issue (some programming language limit recursion depth).

Note, great care must be taken in interpreting these results. Our intention was to create as realistic graphs as possible and then show that the relative performance of PReview greatly outperformed that of the Medidata and NIST PM solutions. In this work, we have done that both theoretically and, in this section, empirically. However, NGAC graphs from operational deployments may have different parameter values or properties not modeled by our graph simulator. Such differences can greatly affect the absolute timing values (as we saw in our experiments on graph of 2000 nodes in changing the parameter values). Thus, we caution the reader to avoid using this work to calculate a precise upper bound on the size of graph that can be processed by any of the three algorithms. That said, the linear nature of PReview should make it suitable for use on any realistic NGAC graph.

5. ACCESS CONTROL VISUALIZATION

These graph algorithms enable access control decisions to be made while simultaneously instantiating multiple access control policies. However, a major question remaining is how to effectively communicate this set of privileges to the users and enable administrator review of a user’s privileges. To this end we have designed an access control visualization approach that meets the following goals:

1. Leverage the access control graph to create a default visualization method for review of user file access (this avoids administrators having to separately maintain some sort of file hierarchy).
2. Abstract away the access control policy details such that the users (or administrators) do not need to understand the policies nor need to know which of the files are covered by which policies.

A visualization approach meeting these goals will provides a default mechanism to enable efficient review of user privileges.

The NIST PM implementation, version 1.5, meets the first goal by leveraging the access control graph. This approach uses the PM itself as a root node in a file hierarchy and then the instantiated access control policies as the second level folders. Clicking on the access control policies enables the user to traverse the object DAG backwards (displaying only oa nodes pertaining to the chosen policy) until reaching the desired files. Unfortunately, this approach does not meet our second goal because their system requires one to navigate to the user’s files by knowing which files are covered by which policies.

Our solution is to use the user’s name as the root in a hierarchical file structure. The second level ‘folders’ are the labels for the border oa nodes reachable from the user’s u node. Given the importance of the border edges in the NGAC access control definition 1, it is natural to use the border oa nodes as the first layer of file organization for the user. When a user clicks on an oa node name, the next level folders that appear are the oa node predecessors in the object DAG for which the user has some privilege. This graph traversal stops when the user reaches object nodes.

In our approach, we abstract away the complexity of the access control graph to make it appear to the user as if they are traversing the usual hierarchical directory structure used by default in all major operating system. In reality, the user is traversing possibly

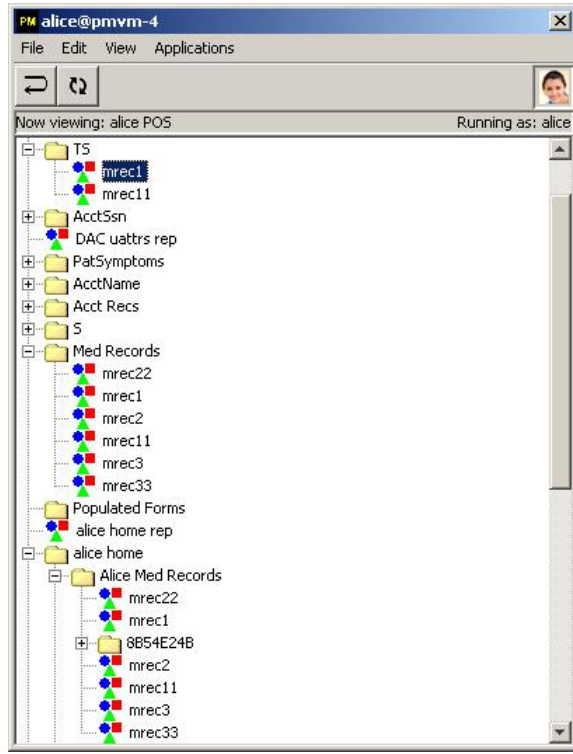


Figure 5: Example hierarchical visualization of a user access rights directed acyclic graph.

overlapping paths of the graph. The number of such paths is exponential and so we perform calculations only on the path actually being traversed by the user. Furthermore, there may be multiple ways for a user to access a particular file. This enables built in flexibility that previously had to be provided explicitly with artifacts such as symbolic links.

Figure 5 provides an example view of a user’s accessible objects taken from one of our testing datasets covering a medical scenario. While it appears to be a typical file hierarchy, note how there are multiple paths by which to traverse to particular files (demonstrating that we are actually traversing a graph). For example, file ‘mrec1’ is available via three different paths in the graph:

1. root \rightarrow TS
2. root \rightarrow Med Records
3. root \rightarrow alice home \rightarrow Alice Med Records.

In fact, all files shown in this visualization depict this multi-path behavior except for the files ‘DAC uattr rep’ and ‘alice home rep’.

5.1 Visualization Algorithm

We now provide an efficient algorithm to determine what files and folders to show when a user clicks on some ‘folder’. Initially, this will be one of the labels for the border oa nodes reachable from the user node, u_1 (which the user can always view by definition). The algorithm is as follows:

1. Let the ‘folder’ on which the user clicks correspond to an oa node, x (note that this algorithm assumes that x is a node that u_1 has the ability to view).

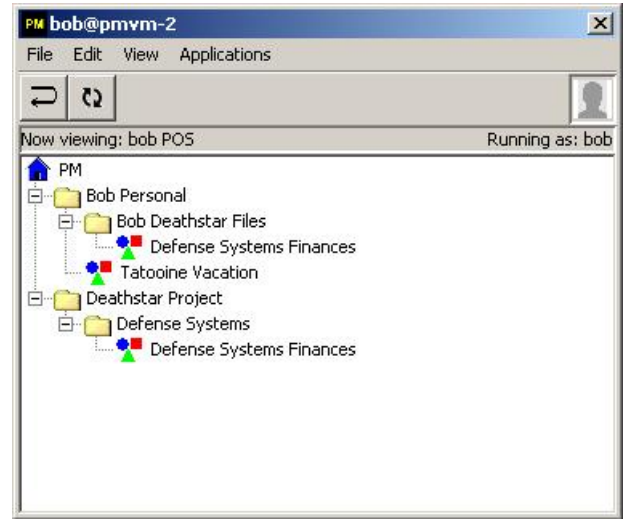


Figure 6: Example File Hierarchy for Access Control Graph in Figure 2.

2. Execute a variant of PReview where in step 4, the algorithm stores all visited nodes instead of just objects when performing the backwards BFS. Thus, we have identified a set of nodes of ‘interest’ that includes both o and oa nodes.
3. In step 5 of the PReview algorithm, use the set of predecessors of x as the input nodes and intersect this set with the set of nodes of interest (from the previous step in PReview) to create a new set of nodes of interest. This is analogous to what already occurs in step 4 except that instead of using a single object as input we are using x ’s set of predecessor nodes.
4. PReview will return whether or not each predecessor of x is visible to u_1 . Display those predecessors that are visible.

This algorithm is a simple variant on PReview and thus is linear, $O(n + m)$ (assuming as usual that the number of distinct access right types and policy classes are a small constant).

5.2 Visualization Examples

We now return to our example of the accountant Bob who is working on the defense systems finances for a deathstar. We will use our visualization approach to show the files available to Bob in Figure 2 using the node labels from Table 1.

The fully available hierarchical tree for user Bob is shown in Figure 6. This assumes that Bob has clicked on the ‘Bob Personal’ folder followed by a click on the ‘Bob Deathstar Files’ subfolder. It also assumes that Bob has clicked on the ‘Deathstar Project’ folder followed by a click on the ‘Defense Systems’ folder. These four clicks expand out visually all of Bob’s available folders and files as shown in Figure 6. Note that the ‘Technical Designs’ folder and the ‘Energy Shield’ file are not visible because they are not accessible to user Bob.

A feature of this approach is that user Bob has access to the ‘Deathstar Finances’ file through both his own documents folder as well as the ‘Deathstar Project’ folder (logically this is because Bob is the owner/maintainer of that file). This again demonstrates the power of the approach where the user visually sees a hierarchy but can access the same files through multiple paths (without the need to explicitly create such linkages).




Figure 7: Minimal Access Control Graph Containing an Orphaned File.

Note that while Bob has access to the ‘Deathstar Project’ folder, he is unable to see anything in the ‘Technical Designs’ folder including the ‘Energy Shield’ file. For Bob to be able to access the ‘Technical Designs’ sub-folder and ‘Energy Shield’ file, there would have to exist an edge $oa_5 \rightarrow oa_2$, $oa_3 \rightarrow oa_2$, $oa_5 \rightarrow oa_1$, or $oa_3 \rightarrow oa_1$ (see Figure 2). Alternately, the existence of an edge $o_3 \rightarrow oa_1$ or $o_3 \rightarrow oa_2$ would be sufficient to allow Bob access to the ‘Energy Shield’ file per Definition 1. However, for this our visualization approach would not allow Bob to use the ‘Technical Designs’ folder because it would still not be accessible to Bob. In this case, Bob could access the ‘Energy Shield’ file through the folder ‘Bob Personal’. Thus, when a user can’t get to one of their files through some particular oa node, there generally exists another oa node that will permit access through the visualization approach.

5.3 Orphan Files

However, there does exist the possibility that a user may not be able to traverse the visualization to reach a file that by Definition 1 is accessible. We call such files ‘orphan’ objects. None of the examples in the NGAC or the PM specification will generate orphans. Likewise, in our own test data sets we have never experienced an orphan file. Nevertheless, the possibility exists and so we discuss approaches to address this eventuality.

For an orphan file to exist for a particular user, there must be an object node that is accessible but each path from the object to the set of reachable border oa nodes has a node that is not accessible. This happens when for each path, there exists an intermediate node that ‘requires’ a policy class not provided by the path’s border oa node (or any other path from that node to a reachable border oa node). Note that an intermediate oa node requires a policy class when it has a path to the corresponding pc node (see Definition 1). While the intermediate nodes on each path are not accessible, each path still provides the user privileges to the object such that the union of the received privileges enables the object to be accessible.

Figure 7 shows the simplest possible access control graph with an orphan file. o_1 is accessible because it receives pc_1 read privileges from oa_2 and pc_2 read privileges from oa_1 . However, oa_3 is not accessible because it requires pc_1 privileges but only receives pc_2 privileges from oa_1 . Likewise, oa_4 is not accessible because it requires pc_2 privileges but only receives pc_1 privileges from oa_2 .

We have identified three different approaches to handling the possibility of orphan files such that the user can still find and access them:

1. Enable the user to perform a search through all accessible files as a method to have access to any orphaned files. Our PReview algorithm from section 4.1 can provide a list of all accessible files, both orphaned and available through the vi-

ualization approach. The user can simply perform a regular expression search on that list.

2. In the user’s visualization of their file hierarchy, provide a folder at the second tier (alongside the reachable border oa node labels) that is labeled ‘Orphan Files’. The orphan files can be detected when first launching the visualization and then listed in that directory. Our linear time algorithm for finding orphan files is provided below.
3. Show the user orphaned files while they are traversing their hierarchical file structure. Whenever a non-accessible folder is encountered, perform a search for orphaned nodes only above the non-accessible folder. If orphans are encountered then show them in the current directory with a special designation to indicate that they are orphans. This can be accomplished through using our linear time algorithm to find orphans (below) in combination with a reverse BFS from the non-accessible folder.

Each of these three approaches can be used independently or together simultaneously.

5.4 Orphan Node Detection Algorithm

This algorithm, based on PReview, enables one to detect orphan nodes for a user node u_1 in linear time. The algorithm is as follows:

1. Execute a variant of PReview where in step 4, the algorithm stores all visited nodes instead of just objects when performing the backwards BFS. Thus, we have identified a set of nodes of ‘interest’ (including both o and oa nodes).
2. PReview will return the o and oa nodes that are accessible.
3. Analogous to steps 1-4 in the PReview algorithm, perform a backwards BFS from the set of reachable border oa nodes. However, this time only visit a node if it is accessible (determined in the previous step). Record this set of visible nodes.
4. The set of accessible nodes (from step 3) that are not visible (from step 4) is the set of orphan nodes.

6. CONCLUSION

It is vital to limit insider access to only the data necessary to perform their work in order to limit possible leakage of sensitive data to outside entities. To best constrain this access, we need scalable access control methodologies that support simultaneous instantiation of multiple access control policies (e.g., DAC and MAC).

The NGAC provides a solution to this important problem by enabling the instantiation of multiple security policies within a single access control system. It quite appropriately provides requirements without specifying implementation details, allowing for competing approaches. However, the existing reference implementations use cubic algorithms, which raised the serious question as to whether or not NGAC is scalable. Furthermore, NGAC did not provide guidance on how to visualize the results of the systems, making it unclear how perform review of user access privileges.

This work addresses both of these issues. In [5] the authors pointed out that performing policy review on ABAC models, such as NGAC, is polynomial time. We provide the first implementation of NGAC using an efficient linear time algorithm (bounded to the parts of the graph relevant to the user). With minimal modifications the methods used in PReview can be adapted for other types of policy review as well. Such as identifying the users that

can access and object. Furthermore, we provide a novel visualization approach that works by default with multiple access control policies and that enables efficient linear time review of user access rights. Given that the only other multi-policy approach with available reference implementations has been shown to not be scalable, our work thus enables for the first time a scalable methodology for limiting user access through simultaneous instantiation of multiple policies.

7. ACKNOWLEDGMENTS

The authors would like to thank Aaron Zinger from Medidata for his help in understanding their implementation. We would also like to thank David Ferraiolo from NIST for his guidance in learning nuances of the NGAC standard and the associated NIST Policy Machine implementation. Lastly, we would like to thank Richard Harang from the United States Army Research for his time in reviewing our paper.

8. REFERENCES

- [1] ANSI. American national standard for information technology, role-based access control (RBAC), 2004.
- [2] ANSI. American national standard for information technology - next generation access control - functional architecture (NGAC-FA), 2013.
- [3] ANSI. American national standard for information technology - next generation access control - generic operations and data structures (NGAC-GOADS), 2016.
- [4] K. Belyaev. tinyPM Prototype. www.github.com/kirillbelyaev/tinypm, 2015.
- [5] P. Biswas, R. Sandhu, and R. Krishnan. Label-based access control: An ABAC model with enumerated authorization policy. In *Proceedings of the 2016 ACM International Workshop on Attribute Based Access Control*, ABAC '16, pages 1–12, New York, NY, USA, 2016. ACM.
- [6] B. Bollobas. *Random graphs*. Cambridge studies in advanced mathematics. Cambridge university press, Cambridge, New York (N. Y.), Melbourne, 2001.
- [7] D. F. Brewer and M. J. Nash. The chinese wall security policy. In *Security and Privacy, 1989. Proceedings., 1989 IEEE Symposium on*, pages 206–214. IEEE, 1989.
- [8] D. Ferraiolo, V. Atluri, and S. Gavrila. The policy machine: A novel architecture and framework for access control policy specification and enforcement. *Journal of Systems Architecture*, 57(4):412–424, 2011.
- [9] D. Ferraiolo, S. Gavrila, and W. Jansen. Policy machine: Features, architecture, and specification. Technical Report NISTIR 7987 Revision 1, National Institute of Standards and of the core services of the central IT organization.
- [10] GitHub. Github code repository. www.github.com, 2016.
- [11] X. Jin, R. Krishnan, and R. Sandhu. *A Unified Attribute-Based Access Control Model Covering DAC, MAC and RBAC*, pages 41–55. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [12] Medidata Solutions Worldwide. Medidata Policy Machine code on github, version 1.1.0., www.github.com/mdsol/the_policy_machine, 2016.
- [13] NCSC. *A Guide to Understanding Discretionary Access Control in Trusted Systems*. Number NCSC-TG-003. National Computer Security Center, Fort George G. Meade, Maryland, USA, 1 edition, Sept. 1987.
- [14] NIST. NIST Policy Machine code on github, version 1.5., www.github.com/PM-Master/PM, 2016.
- [15] OASIS. eXtensible access control markup language (XACML) Version 3.0., OASIS Standard, Jan. 2013.
- [16] Organization for the advancement of structured information standards OASIS. www.oasis-open.org, 2016.
- [17] D. Servos and S. L. Osborn. *HGABAC: Towards a Formal Model of Hierarchical Attribute-Based Access Control*, pages 187–204. Springer International Publishing, Cham, 2015.
- [18] F. Turkmen and B. Crispo. Performance evaluation of XACML PDP implementations. In *Proceedings of the 2008 ACM Workshop on Secure Web Services*, SWS '08, pages 37–44, New York, NY, USA, 2008. ACM.
- [19] U.S. Department of Defense. Trusted computer system evaluation criteria DoD 5200.28-STD, 1985.
- [20] E. Yuan and J. Tong. Attributed based access control (ABAC) for web services. In *IEEE International Conference on Web Services (ICWS'05)*, page 569, July 2005.